Real-time Collaborative Environments and the Grid

Colin Perkins
University of Glasgow

Ladan Gharai
USC Information Sciences Institute

Information Sciences Institute

UNIVERSITY of GLASGOW

Overview

- Real-time Collaborative Environments
  - Vision
  - Progress and Limitations
- Building the Supporting Infrastructure
- Proposal for a New Architecture
- Open Issues and Related Work

Real-Time Collaborative Environments

Original Picture ©1998 UCL

- Support collaborative work by scientists, industry and others
  - Group communication scenarios
  - Widely distributed participants
- Integrate with Grid computing resources
  - Visualization facilities
  - Shared state and data repositories, computational resources, applications
- Provide a high-quality audio-visual experience
  - High fidelity audio, possibly with spatial positioning
  - High resolution video, wide colour gamut, high frame rate
- Provide a sense of community and presence
  - Venues as a rendezvous point; virtual meeting place
- Secure and trusted infrastructure

Progress and Limitations

- Networked teleconferencing environments have become flexible, secure, and integrated with Grid computing environments
- Standards have matured; commercial teleconferencing and VoIP
- Facilitated by greatly increased network and host performance

- Media quality has stagnated
  - Tools are more readily available
  - But: media capabilities are all-b
    and don't leverage changes in te

Starting to be addressed in AccessGrid:
• DV from Australian National University
• DV + compressed HD from GIST, Korea
• Uncompressed HD from USC/ISI
ConferenceXP

- The network has become fragmented
  - Multicast still not widely available
  - Growing prevalence of NAT and firewalls

Open issue - focus of this talk

Building the Supporting Infrastructure

How are collaborative environments built?

- Many to many collaborative work
- Audio/Video + shared application
- Predictable virtual environment with secure and persistent state

Multicast protocols for scalable data distribution

Venue server acts as rendezvous point, encourages communities, enforces security & keeps state

Mbone tools + SDR
AccessGrid
ConferenceXP

Infrastructure Challenges

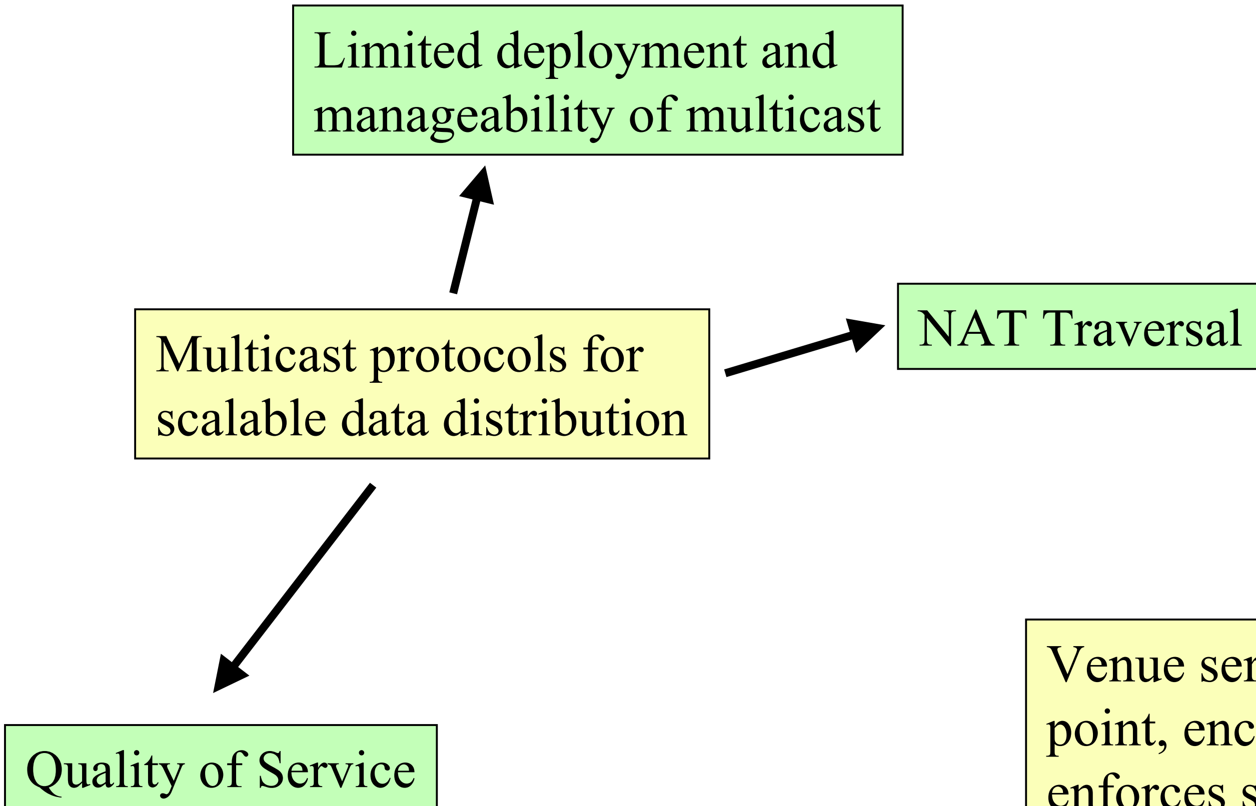

Venue server acts as rendezvous point, encourages communities, enforces security & keeps state

Multicast Deployment

- How to support…
  - Groups with 10s (100s?) of participants
  - Rich many-to-many interaction between participants
  - Varying participant roles (audience; speaker; chair; question)
  - Control of participation, rights and responsibilities
- Want to use IP multicast, but:
  - Difficult to deploy and unstable
    - Local issues with old Ethernet switches flooding traffic
    - PIM+MSDP suffers from interdomain flooding of source state, unstable
    - Many security issues, easy to mount DoS attacks
  - Difficult to manage
    - Traffic engineering not well defined, limited QoS monitoring
    - Infrastructure not stable

⇒ Strong push within IETF to deprecate traditional IP multicast in favour of source-specific (one-to-many) multicast

Multicast Deployment

What are the scalable alternative to multicast?

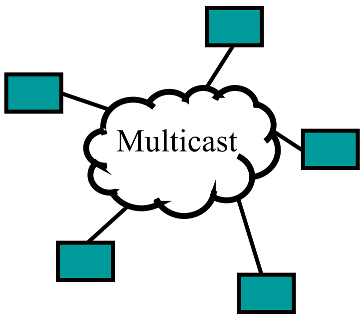

Hard to deploy due to problems with multicast

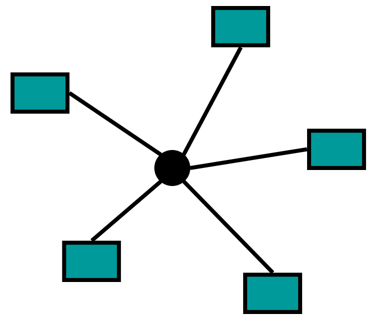

Too much load on server; central point of failure

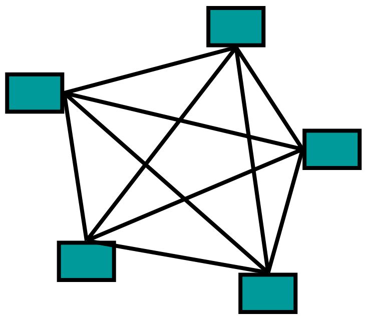

Distributes load; resilient to failures

⇒ Peer-to-peer overlays or application level multicast, independent of the network layer

NAT Traversal

- Growing prevalence of Network Address Translation (NAT) and Firewalls is fragmenting the Internet
  - Complicates applications since they cannot easily name/access peers
  - Hosts no longer have unique addresses
  - Bidirectional connectivity not assured, may vary by protocol or direction
  - Especially affects protocols with dynamic signalling ⇒ collaborative tools
- Numerous solutions ("kludges") under development in IETF:
  - STUN
  - TURN
  - ICE

  Methodologies for detecting presence of NAT, deducing it's behaviour, and establishing connectivity
- Signalling driven NAT detection and connection establishment

⇒ Leverage the venue server as a signalling proxy; reuse protocols

Infrastructure Challenges + Solutions

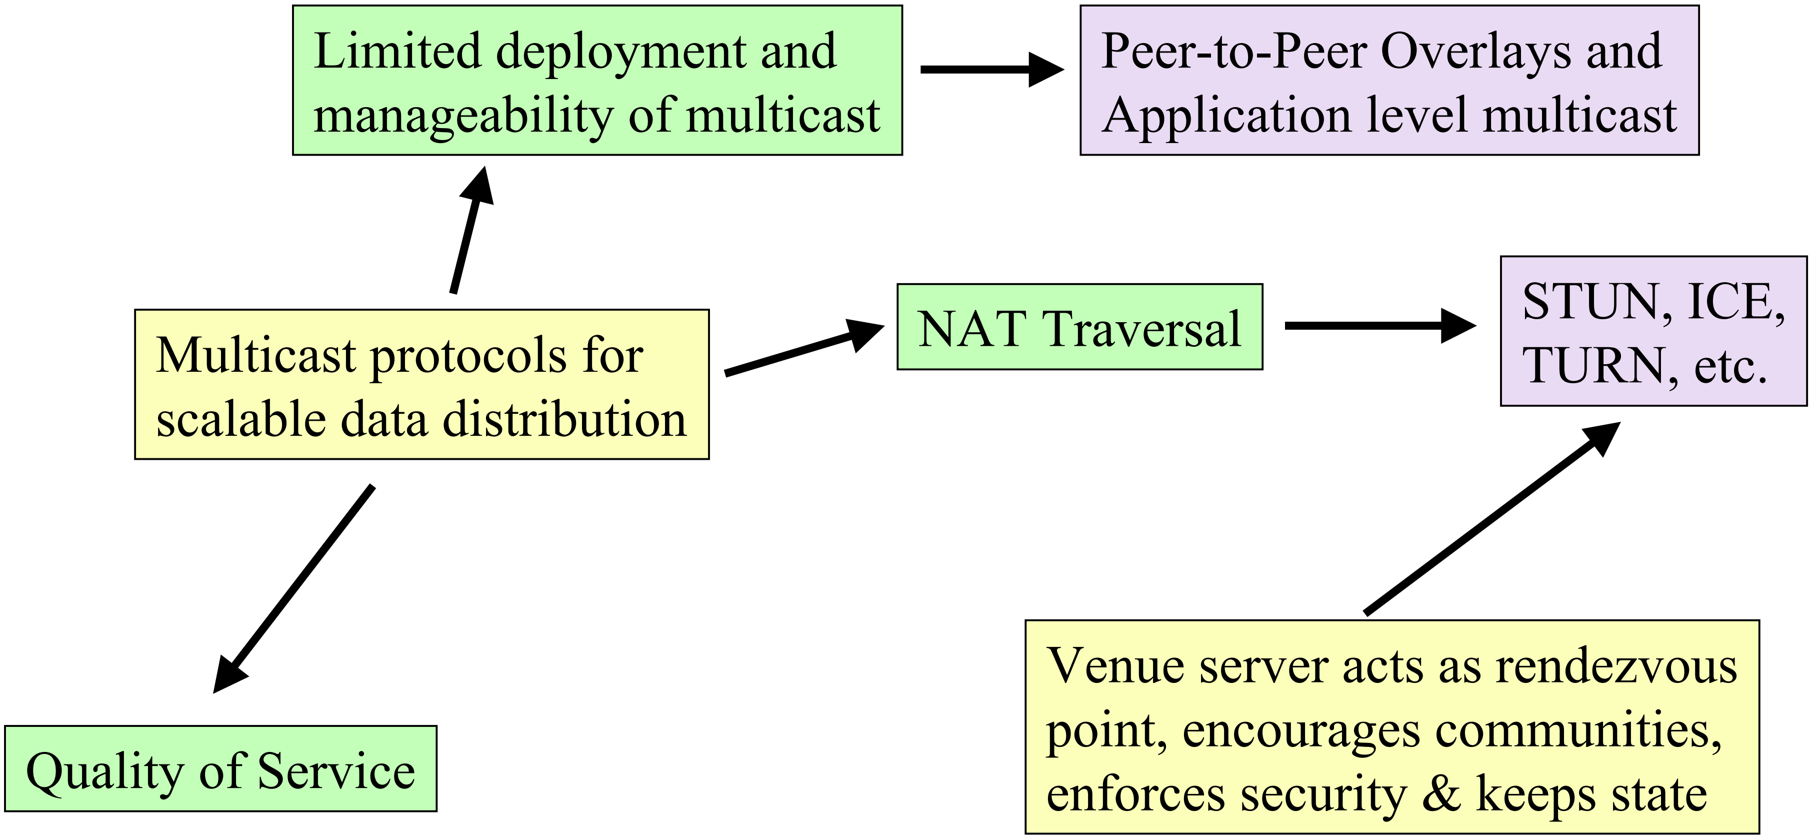

$\Rightarrow$ a peer-to-peer overlay with NAT traversal

A New Architectural Direction?

- Established that an overlay with NAT traversal is needed
- Desirable that this…
  - be implemented as reusable middleware
  - be application independent
  - allows existing applications to run unchanged
  - be simple for network administrators who must deploy it and enforce security policy

⇒ Propose middleware to build an overlay IP network, with NAT traversal and multicast support, run applications on that

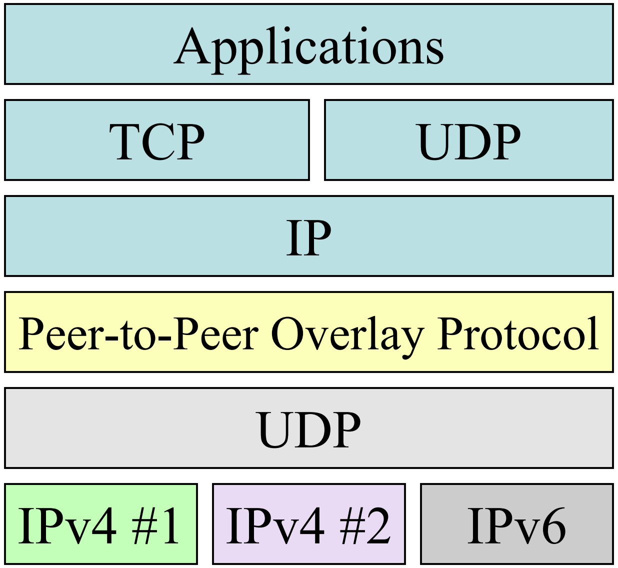

Outline System Operation

1. Rendezvous
   – Participants rendezvous at venue server
   – Authenticate and acquire venue state
   – Venue server acts as signalling relay; may act as last-ditch media relay
2. Session Initiation Signalling
   – Participants submit list of their possible network addresses to venue server
   – Venue server coordinates pair-wise ICE exchange between participants, as a UDP-based connectivity check
3. Building the overlay
   – Build IP overlay using tunnelled UDP/IP on address/port pairs determined by the connectivity check
   – Any appropriate peer-to-peer overlay building protocol
4. Running Collaborative Work Tools
   – Start media tools using IP addresses on the overlay network

An IP Overlay Network on UDP/IP

Advantages

- Simple, well-known, API for user applications
- Runs existing tools; reusable
- Single address space, irrespective of underlying addressing realms
- Doesn't require multicast in the underlying network; easy to support multicast in the overlay
- Overlay has restricted, and known, membership

Disadvantages

- Bandwidth and processing overhead due to layering and tunnelling
- Complexity of implementation
  - Although much of this needed for any system in today's Internet
- The "yeuch!" factor

Discussion Items and Open Issues

- Complexity of signalling
  - Unavoidable once you accept NAT/firewall + limited native multicast
  - No worse than similar solutions for telephony
  - Can be bundled into a middleware solution
- Quality of service support not yet well defined
  - NAT traversal requires a connectivity check
  - Potentially possible to reuse this to check non-IP connectivity
    - E.g. Try to negotiate an MPLS path or RSVP reservation
  - For future study…

Summary + Related Work

- Real-time collaborative environments will need to deal with:
  - Fragmented connectivity, due to NAT/Firewall
  - Lack of native multicast support
  - Lack of QoS support
- Existing systems fail at this
- Many ad-hoc solutions for certain applications
- Propose to encapsulate these in a middleware component
  - An IP network built on a UDP/IP-based overlay
  - Restores end-to-end multicast connectivity for applications

- Related work: a synthesis of…
  - Microsoft peer-to-peer SDK with Teredo
  - Xbone
  - Peer-to-peer protocols